arXiv:1607.06023v1 [cs.NI] 28 Jun 2016

Modeling wireless network routing using sheaves

Michael Robinson
Mathematics and Statistics
American University
Washington, DC, USA
michaelr@american.edu

July 26, 2016

Abstract

This article explains how to construct a sheaf model for passing traffic through a wireless network with a single channel carrier sense multiple access/collision detection (CSMA/CD) media access model.

1 Introduction

We make the following *single channel assumption*: if a link connected to a node is jammed, then that node cannot receive transmissions from *any* other node.

Definition 1. An *abstract simplicial complex* X on a set A is a collection of ordered subsets of A that is closed under the operation of taking subsets. We call an element of X which itself contains $k+1$ elements a *k-cell*. We usually call a 0-cell a *vertex* and a 1-cell an *edge*.

If a, b are cells with $a \subset b$, we say that a is a *face* of b, and that b is a *coface* of a. A cell of X that has no cofaces is called a *facet*.

The *closure* $\operatorname{cl} Y$ of a set Y of cells in X is the smallest abstract simplicial complex that contains Y. The *star* $\operatorname{star} Y$ of a set Y of cells in X is the set of all cells that have at least one face in Y.

Suppose a radio network consisting of a collection of nodes $N = \{n_i\}$ is active in a spatial region R. Assume all nodes communicate through a single-channel, broadcast resource. An open set $U_i \subset R$ is associated to each node n_i that represents its *transmitter coverage region*. For each node n_i, a continuous function $s_i : U_i \to \mathbb{R}$ represents its *signal level* at each point in U_i. Without loss of generality, we assume that there is a global threshold T for accurately decoding the transmission from any node.

Definition 2. The *link graph* is the following collection of subsets of N:

1. $\{n_i\} \in N$ for each node n_i, and

2. $\{n_i, n_j\} \in N$ if $s_i(n_j) > T$ and $s_j(n_i) > T$.

The *link complex* $L = L(N, U, s, T)$ is the clique complex of the link graph, which means that it contains all elements of the form $\{i_1, \ldots, i_n\}$ whenever this set is a clique in the link graph.

Proposition 3. *Each facet in the link complex is a maximal set of nodes that can communicate directly with one another (with only one transmitting at a time).*

Proof. Let c be a cell of the link complex. By definition, for each pair of nodes, $i, j \in c$ implies that $s_i(n_j) > T$ and $s_j(n_i) > T$. Therefore, i and j can communicate with one another. □

Corollary 4. *Facets of the link complexes represent common broadcast resources.*

2 Interference from a transmission

The interference caused by a transmission impacts the usability of the network outside of the transmission's immediate vicinity. This section builds a consistent definition of the *region of influence* of a node or a link within the network. To justify this definition, we use a local model that describes which configurations of nodes can transmit simultaneously.

Definition 5. Suppose that X is a simplicial complex (such as an interference or link complex) whose set of vertices is N. Consider the following assignment $\mathcal{A}$ of additional information to capture which nodes are transmitting and decodable:

1. To each cell $c \in X$, assign the set

$$\begin{aligned} \mathcal{A}(c) \quad = \quad & \{n \in N : \text{there exists a cell } d \in X \text{ with} \\ & c \subset d \text{ and } n \in d\} \cup \{\perp\} \end{aligned}$$

   of nodes that have a coface in common with c, along with the symbol $\perp$. We call $\mathcal{A}(c)$ the *stalk* of $\mathcal{A}$ at c.

2. To each pair $c \subset d$ of cells, assign the *restriction function*

$$\mathcal{A}(c \subset d)(n) = \begin{cases} n & \text{if } n \in \mathcal{A}(d) \\ \perp & \text{otherwise} \end{cases}$$

For instance, if $c \in X$ is a cell of a link complex, $\mathcal{A}(c)$ specifies which nearby node is transmitting and decodable, or $\perp$ if none are. The restriction functions relate the decodable transmitting nodes at the nodes to which nodes are decodable along an attached wireless link. Similarly, if $c \in X$ is a cell of an interference complex, $\mathcal{A}(c)$ also specifies which nearby node is transmitting, and effectively locks out any interfering transmissions from other nodes.

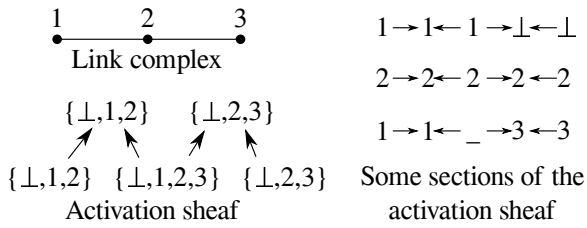

Figure 1: A link complex (left top), sheaf $\mathcal{A}$ (left bottom), and three sections (right). The restrictions are shown with arrows. global section when node 1 transmits (right top), global section when node 2 transmits (right middle), and a local section with nodes 1 and 3 attempting to transmit, interfering at node 2 (right bottom)

Definition 6. The assignment $\mathcal{A}$ is called the *activation sheaf* and is an example of a *cellular sheaf* – a mathematical object that stores local data. The theory of sheaves explains how to extract consistent information, which in the case of networks consists of nodes that whose transmissions do not interfere with one another.

A *section* of $\mathcal{A}$ supported on a subset $Y \subseteq X$ is an assignment $s : Y \to N$ so that for each $c \subset d$ in Y, $s(c) \in \mathcal{A}(c)$ and $\mathcal{A}(c \subset d)\,(s(c)) = s(d)$. A section supported on X is called a *global section*.

Specifically, global sections are complete lists of nodes that can be transmitting without interference.

Example 7. Figure 1 shows a network with three nodes, labeled 1, 2, and 3. When node 1 transmits, node 2 receives. Because node 2 is busy, its link to node 3 must remain inactive (right top). When node 2 transmits, both nodes 1 and 3 receive (right middle). The right bottom diagram shows a local section that cannot be extended to the cell marked with a blank. This corresponds to the situation where nodes 1 and 3 attempt to transmit but instead cause interference at node 2.

Definition 8. Suppose that s is a global section of $\mathcal{A}$. The *active region* associated to a node $n \in X$ in s is the set

$$\text{active}(s, n) = \{a \in X : s(a) = n\},$$

which is the set of all nodes that are currently waiting on n to finish transmitting.

Lemma 9. *The active region of a node is a connected, closed subcomplex of X that contains n.*

Proof. Consider a cell $c \in \text{active}(s, n)$. If c is not a vertex, then there exists a $b \subset c$; we must show that $b \in \text{active}(s, n)$. Since s is a global section $\mathcal{A}(b \subset$

$c)s(b) = s(c) = n$. Because $s(c) \neq \perp$, the definition of the restriction function $\mathcal{A}(b \subset c)$ implies that $s(b) = n$. Thus $b \in \text{active}(s, n)$ so $\text{active}(s, n)$ is closed.

If $c \in \text{active}(s, n)$, then c and n have a coface d in common. Since s is a global section $s(d) = \mathcal{A}(c \subset d)s(c) = \mathcal{A}(c \subset d)n = n$. Thus, $n \in \text{active}(s, n)$, because n is a face of d and $\text{active}(s, n)$ is closed. This also shows that every cell in $\text{active}(s, n)$ is connected to n. □

Lemma 10. *The star over the active region of a node does not intersect the active region of any other node.*

Proof. Let $c \in \text{star active}(s, n)$. Without loss of generality, assume that $c \notin \text{active}(s, n)$. Therefore, there is a $b \in \text{active}(s, n)$ with $b \subset c$. By the definition of the restriction function $\mathcal{A}(b \subset c)$, the assumption that $c \notin \text{active}(s, n)$, and the fact that s is a global section, $s(c)$ must be $\perp$. □

Corollary 11. *If s is a global section of an activation sheaf $\mathcal{A}$, then the* support *of s – the set of cells c where $s(c) \neq \perp$ – consists of a disjoint union of active regions of nodes.*

Lemma 12. *The active region of a node is independent of the global section. More precisely, if r and s are global sections of $\mathcal{A}$ and the active regions associated to $n \in X$ are nonempty in both, then active(s, n) = active(r, n).*

Proof. Without loss of generality, we need only show that $\text{active}(s, n) \subseteq \text{active}(r, n)$. If $c \in \text{active}(s, n)$, there must be a cell $d \in X$ that has both n and c as faces. Now $s(n) = r(n) = n$ by Lemma 9, which means that $r(d) = \mathcal{A}(n \subset d)r(n) = n$. Therefore, since $\text{active}(r, n)$ is closed, this implies that $c \in \text{active}(r, n)$. □

Corollary 13. *The space of global sections of an activation sheaf consists of all sets of nodes that can be transmitting simultaneously without interference.*

Definition 14. Because of the Lemmas, we call the star over an active region associated to a node n the *region of influence*. The region of influence of a node can be written as a union

$$\text{roi } n = \bigcup_{f \in F} \text{star cl } f,$$

for a collection F of facets, so we call the star over a closure of a facet the *region of influence* of that facet.

Corollary 15. *The complement of the region of influence of a facet is a closed subcomplex.*

Although the space of global sections for an activation sheaf is a useful invariant, the cohomology of an activation sheaf is rather uninteresting. However, activation sheaves are not sheaves of vector spaces, so we need to enrich their structure somewhat to see this.

Definition 16. If $\mathcal{A}$ is an activation sheaf on an abstract simplicial complex X, the *vector activation sheaf* $\widehat{\mathcal{A}}$ is given by specifying its stalks and restrictions:

1. To each cell $c \in X$, let $\widehat{\mathcal{A}}(c)$ be the vector space whose basis is $\mathcal{A}\backslash\{\perp\}$ (so the dimension of this vector space is the cardinality of $\mathcal{A}$ without counting $\perp$)

2. The restriction map $\widehat{\mathcal{A}}(c \subset d)(n)$ is the basis projection, which is well-defined since $\mathcal{A}(d) \subseteq \mathcal{A}(c)$.

Theorem 17. *The dimension of the cohomology spaces of a vector activation sheaf $\widehat{\mathcal{A}}$ on a link complex X are*

$$\dim H^k(\widehat{\mathcal{A}}) = \begin{cases} \text{the total number of nodes} & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases}$$

Proof. Every global section of $\mathcal{A}$ corresponds to a global section of $\widehat{\mathcal{A}}$, but formal linear combinations of global sections of $\mathcal{A}$ are also global sections of $\widehat{\mathcal{A}}$. Therefore, a global section of $\widehat{\mathcal{A}}$ merely consists of a list of those nodes that are transmitting, without regard for whether they interfere.

The fact that the other cohomology spaces are trivial is considerably more subtle. Consider the decomposition

$$X = \bigcup_i F_i$$

of the link complex into the set of its facets. Suppose that F_i is a facet of dimension k, and defining $\mathcal{F}_i$ to be the direct sum of $k+1$ copies of the constant sheaf supported on F_i. (Each copy corresponds one of the vertices of F_i.) Then there is an exact sequence of sheaves

$$0 \to \widehat{\mathcal{A}} \xrightarrow{\Delta} \bigoplus_i \mathcal{F}_i \xrightarrow{m} \mathcal{S} \to 0$$

where Δ is a map that takes a basis vector corresponding to a given node to the linear combination of all corresponding basis vectors in each copy of the constant sheaves, and m is therefore a kind of difference map. This exact sequence leads to a long exact sequence

$$\cdots H^{k-1}(\mathcal{S}) \to H^k(\widehat{\mathcal{A}}) \to \bigoplus_i H^k(\mathcal{F}_i) \to H^k(\mathcal{S}) \cdots$$

Since each $\mathcal{F}_i$ is a direct sum of constant sheaves supported on a closed subcomplex, it only has nontrivial cohomology in degree 0.

Observe that $\mathcal{S}$ is a sheaf supported on sets of cells lying in the intersections of facets. By Corollary 15, $\mathcal{S}$ must be a direct sum of copies of constant sheaves supported on closed subcomplexes, like each $\mathcal{F}_i$. Thus $\mathcal{S}$ only has nontrivial cohomology in degree 0, which means that for $k > 1$, $H^k(\widehat{\mathcal{A}}) = 0$.

It therefore remains to address the $k = 1$ case, which comes about from the exact sequence

$$\bigoplus_i H^0(\mathcal{F}_i) \to H^0(\mathcal{S}) \to H^1(\widehat{\mathcal{S}}) \to 0.$$

The leftmost map is surjective, since every global section of $\mathcal{S}$ is given by specifying a single transmitting node. By picking exactly one facet containing that node, a global section of the corresponding $\mathcal{F}_i$ may be selected in the preimage. Thus the map $H^0(\mathcal{S}) \to H^1(\widehat{\mathcal{S}})$ must be the zero map and yet also surjective. This completes the proof. □

3 Modeling link and node data payloads

The activation sheaf describes the state of the network at a single instant in time. Because the network conditions may change over time, the link and interference complexes may also change with time. This section describes a general framework for representing both these changes and the data that is transmitted over the links.

In order to capture changes in the network's topology over time, it is appropriate to use a single link or interference complex to represent a single timeslice. To represent how the network's state evolves over several consecutive timeslices, we construct additional links between nodes in different timeslices. These links carry information from one timeslice to the next.

Extending the definition of a link complex above, again suppose that a radio network consists of a collection of nodes $N = \{n_i\}$ in a spatial region R in which a coverage region $U_i \subset R$ is associated to each node n_i. For each node n_i, assign a signal level function $s_i : U_i \times \mathbb{Z} \to \mathbb{R}$, where the second input represents time. Again, without loss of generality, we assume that there is a global decoding threshold T.

Definition 18. The *time-dependent link graph* is the following collection of subsets of $N \times \mathbb{Z}$:

1. $\{(n_i, t)\} \in N$ for each node n_i and $t \in \mathbb{Z}$,
2. $\{(n_i, t), (n_j, t)\} \in N$ if $s_i(n_j, t) > T$ and $s_j(n_i, t) > T$, and
3. $\{(n_i, t), (n_i, t+1)\}$ for each node n_i.

The *time-depdendent link complex* $L = L(N, U, s, T)$ is the clique complex of the time-dependent link graph, which means that it contains all elements of the form $\{i_1, \ldots, i_n\}$ whenever this set is a clique in the link graph. The *time t timeslice* of L is the maximal subcomplex of L containing vertices from $N \times \{t\}$. The *time-dependent interference complex* and its timeslices can be defined in an analogous manner.

Figure 2 shows an example of a time-dependent link complex. Notice that each timeslice is a link complex, and that timeslices are attached to one another only by edges between consecutive copies of the same node. The interpretation is that this represents a network in which the links are memoryless; only nodes can retain information from one timeslice to the next.

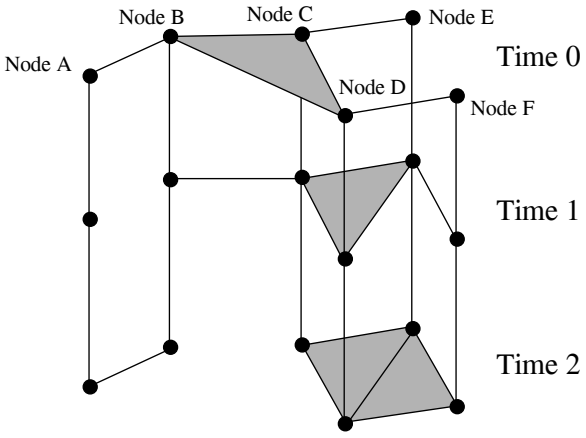

Figure 2: Evolution of a simplicial complex model of a wireless network through time. The nodes A–F listed at time 0 are repeated through the other timeslices as proceeding vertically downward through the diagram.

In order to represent the information that is retained by a node concretely, we will first construct a sheaf model of a queue and then generalize. For example, consider the following diagram of sets

$$\cdots \longrightarrow 0 \longleftarrow A \longrightarrow 0 \longleftarrow A \longrightarrow \cdots$$

in which the arrows represent functions taking all elements of A to 0. This diagram is a sheaf over a simplicial complex whose sections are given by assigning (possibly different) elements of A to each vertex. In other words, the space of sections of this sheaf is the space of A-valued sequences. This sheaf should be interpreted as a 1-deep queue: each A in the diagram above represents the contents of the queue at a timeslice, but nothing is retained from one timeslice to the next. To generalize this construction to larger queues, consider the following definition.

Definition 19. Suppose X is the abstract simplicial complex model of $\mathbb{R}$ whose vertices are given by the set of integers and whose edges are given by pairs $(t, t+1)$. The *n-term grouping sheaf* $\mathcal{A}^{(n)}$ is given by the diagram written over X

$$\longrightarrow A^{n-1} \underset{\sigma_-}{\longleftarrow} A^n \xrightarrow{\sigma_+} A^{n-1} \underset{\sigma_-}{\longleftarrow} A^n \xrightarrow{\sigma_+}$$

where $\sigma_\pm$ are the projection functions given by

$$\sigma_-(x_1, \ldots, x_n) = (x_2, \ldots, x_n)$$

and

$$\sigma_+(x_1, \ldots, x_n) = (x_1, \ldots, x_{n-1}).$$

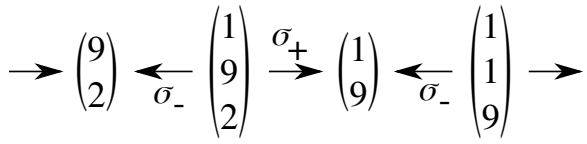

Figure 3: Section of a 3-term grouping sheaf modeling a 3-deep queue

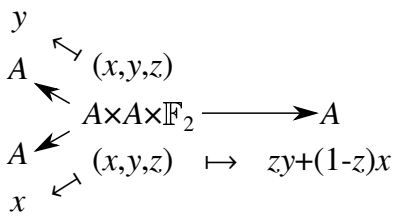

Figure 4: A sheaf that models switching an output (right) between two inputs (left)

The $\sigma_\pm$ projection functions in a grouping sheaf represent the action of advancing the queue from one timestamp to the next. Specifically, consider the diagram shown in Figure 3, which represents a 3-term grouping sheaf. To read the diagram as a 3-deep queue, think of data as entering the queue from the first element of the vectors (the top) and exiting from the last (the bottom). The action of σ_- is simply to drop the first element, which maintains consistency between the current timestep and the one previous. Conversely, σ_+ drops the last element, which links the current timestep to the next one.

In addition to retaining a queue of packets, nodes also need to be able to switch between transmit, receive, and idle states. These states control two behaviors (1) the link activity and (2) when the queues advance. (For the moment, let us only consider the transmit queue in a node.) The simplest way to manage several models of behavior is by using a state variable to switch between them. This is well established in the mathematical systems theory literature [1, 2], though we merely leverage the ideas as needed here without going into more detail.

Consider the sheaf model shown in Figure 4, which exhibits switching behavior. This sheaf represents two inputs (on the left) valued in a set A linking to a single A-valued output (on the right). The data over the vertex contains a boolean state value and both inputs. The left two restriction maps merely extract the two inputs. The right restriction map is given by the expression $zy + (1 - z)x$. Observe that when $z = 0$, this map produces x. On the other hand, if $z = 1$ the map produces y.

The example in Figure 4 is easily generalized to use more elaborate state variables and to exhibit more complex behavior. Consider a wireless network with two nodes $N = \{0, 1\}$ over the course of two consecutive timeslices. The

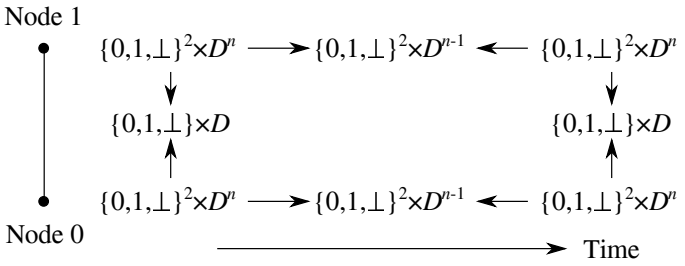

Figure 5: A sheaf representing two timesteps of two nodes with a shared channel

$\{0,1,\perp\}^2\times D^n \longrightarrow \{0,1,\perp\}\times D^{n-1} \longleftarrow \{0,1,\perp\}^2\times D^n$

$(?,0,d_1,...,d_n)$	$\mapsto$	$(0,d_{p(1)},...,d_{p(n-1)})$	Queues RX'ed packet if needed
$(?,1,d_1,...,d_n)$	$\mapsto$	$(1,d_2,...,d_{n-1},0)$	Advances TX queue
$(?,\perp,d_1,...,d_n)$	$\mapsto$	$(\perp,d_2,...,d_{n-1},d_n)$	
		$(0,d_2,...,d_{n-1},d_n)$	$\mapsto (0,?,d_1,...,d_n)$
Advances TX queue		$(1,d_3,...,d_n,0)$	$\mapsto (1,?,d_1,...,d_n)$
		$(\perp,d_2,...,d_{n-1},d_n)$	$\mapsto (\perp,?,d_1,...,d_n)$

Figure 6: The between-timestep restrictions for the sheaf in Figure 5

sheaf we will construct is summarized in Figure 5, which is written over a time-dependent link complex. It is helpful to examine the horizontal restriction maps (within a single node) separately from the vertical restrictions (within a single timeslice).

The diagram shown in Figure 6 represents the state of a single node in two timesteps. The node contains two state variables, a receive buffer, and a transmit queue. The state variables take values in $N \cup \{\perp\}$ where N is the set of nodes in the network, which represents the identity of the node which is currently transmitting or $\perp$ if the link is idle. The state variables lie in a 2-deep queue, so the stalk at a vertex contains both the current and the previous state.

The receive buffer and transmit queue are encoded in the vector D^n: the receive buffer is the first element, and the transmit queue is $n-1$ elements long. Unlike a grouping sheaf, which always advances from one timestep to the next, the transmit queue is only advanced when the node is transmitting. Thus, the stalk over the edge in Figure 6 needs to accomodate a vector that is the same length as the transmit queue to retain the entire queue if the node is not transmitting. When tranmission occurs, the final element of the vector over the edge is set to zero, effectively yielding a grouping sheaf construction.

When the node is receiving, it must make a determination about whether to forward the received packet or not. This decision needs to take into account at least the received packet's destination and priority.

Definition 20. To represent this decision process abstractly – without committing to a specific protocol – this should be modeled as a *receive queue function* $q : D^n \to D^{n-1}$ of the form

$$q(x_1, x_2, \ldots, x_n) = \left(x_{p(1,x_1,\ldots,x_n)}, x_{p(2,x_1,\ldots,x_n)}, \ldots, x_{p(n-1,x_1,\ldots,x_n)}\right)$$

where for brevity, $p(i, x_1, \ldots, x_n)$ is written $p(i)$ in Figure 6.

Notice that a receive queue function q is a permutation of a subset of its inputs that depends on what those inputs are. Different choices for q determine different protocols for forwarding packets, which usually make decisions based on the contents of x_1. For instance, the following are some possible receive queue functions of a node $n_i \in N$ (this list should not be considered exhaustive, nor are any of these protocols optimal in any useful sense)

Forward nothing:

$$q(x_1, \ldots, x_n) = (x_2, x_3, \ldots, x_n).$$

This protocol doesn't forward any received packets – it is assumed that the transmit queues are filled internally by the node.

Forward everything:

$$q(x_1, \ldots, x_n) = (x_1, x_3, \ldots, x_n).$$

This protocol places every received packet at the start of the transmit queue, overwriting whatever happened to be there. Because of this, the protocol may drop packets at the start of the transmit queue. It is helpful to assume that $x_i = 0$ means that nothing is waiting in that slot of the queue.

Forward everything, with queue management:

$$q(x_1, \ldots, x_n) = (0, \ldots, 0, x_1, \ldots, x_n),$$

where x_1 is placed in the last empty slot in the queue so that it gets transmitted sooner than in the previous protocol. If there are no nonzero slots in the queue, then the starting queue entry is dropped and replaced with x_1.

Forward everything for other recipients:

$$q(x_1, \ldots, x_n) = \begin{cases} (0, \ldots, 0, x_1, \ldots, x_n) & \text{if destination of } x_1 \text{ is not } n_i \\ (x_2, x_3, \ldots, x_n) & \text{if destination of } x_1 \text{ is } n_i \end{cases}$$

This protocol forwards every packet that's not destined for the node that received it, by placing it at the start of the queue. Packets that are destined for the receiving node are not forwarded, so it is effectively a combination of the previous two protocols. Again, the protocol may drop packets from the start of the transmit queue.

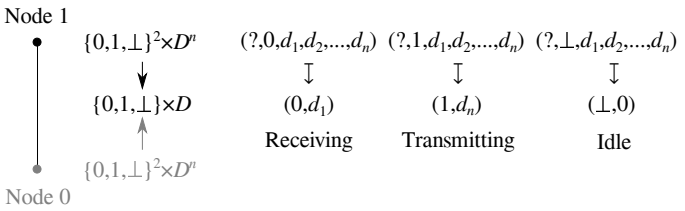

Figure 7: The within-timestep restrictions for the sheaf in Figure 5

Forward everything for other recipients (two priority levels):

$$q(x_1,\ldots,x_n) = \begin{cases} (x_3,\ldots,x_n,x_1) & \text{if destination of } x_1 \text{ is not } n_i, \\ & \text{and priority of } x_1 \text{ is HIGH} \\ (0,\ldots,0,x_1,\ldots,x_n) & \text{if destination of } x_1 \text{ is not } n_i, \\ & \text{and priority of } x_1 \text{ is LOW} \\ (x_2,x_3,\ldots,x_n) & \text{if destination of } x_1 \text{ is } n_i \end{cases}$$

This protocol is the same as the previous one, however packets marked as HIGH priority jump to the end of the transmit queue (next to transmit), dropping any packet at the start of the transmit queue.

Within a timestep, the restriction maps are based on those of an activation sheaf. As Figure 7 shows, the current state variable determines not only the identity of the node utilizing the link (if any) but also the data present on that link. Observe that when a node is transmitting, the last element of the transmit buffer is what appears on the link. When the node is receiving, the link's contents match with the receive buffer.

Using the above example as a recipe, it is straightforward to define a similar sheaf over a time-dependent link complex.

Definition 21. A *data payload sheaf* $\mathcal{D}$ over a time-dependent link complex L with nodes N is parameterized by

1. A vector space D of possible packets, and
2. A transmit queue length $n-1$.

The stalks of $\mathcal{D}$ are given by

For each vertex c of L: $(\{a \in N : \text{there exists a cell } d \in X \text{ with } c \subset d \text{ and } a \in d\} \cup \{\perp\})^2 \times D^n$

For each edge of the form $((c,t),(c,t+1))$: $(\{a \in N : \text{there exists a cell } d \in X \text{ with } c \subset d \text{ and } a \in d\} \cup \{\perp\}) \times D^{n-1}$

For all other simplices c of L: $(\{a \in N : \text{there exists a cell } d \in X \text{ with } c \subset d \text{ and } a \in d\} \cup \{\perp\}) \times D$

The restrictions of $\mathcal{D}$ are given by

1. Between timeslices

$$\mathcal{D}\left((a,t) \subset ((a,t),(a,t+1))\right)(n_1, n_2, x_1, \ldots, x_n) = \begin{cases} (n_2, x_2, \ldots, x_{n-1}, 0) & \text{if } n_2 = a \text{ and } x_n \neq 0 \\ (n_2, x_{p(1)}, \ldots, x_{p(n-1)}) & \text{if } n_2 \neq a \text{ and } n_2 \neq \perp \\ (\perp, x_2, \ldots, x_n) & \text{otherwise} \end{cases}$$

where $(x_1, \ldots, x_n) \mapsto (x_{p(1)}, \ldots, x_{p(n-1)})$ is a receive queue function.

$$\mathcal{D}\left((a,t+1) \subset ((a,t),(a,t+1))\right)(n_1, n_2, x_1, \ldots, x_n) = \begin{cases} (n_1, x_3, \ldots, x_n, 0) & \text{if } n_1 = a \\ (n_1, x_2, \ldots, x_n) & \text{if } n_1 \neq a \text{ and } n_1 \neq \perp \\ (\perp, x_2, \ldots, x_n) & \text{otherwise} \end{cases}$$

2. Within timeslices, all restrictions between simplices $a \subset b$ of dimension 1 or higher are of the form

$$\mathcal{D}\left((a,t) \subset (b,t))\right)(n,x) = \begin{cases} (n,x) & \text{if } (n,x) \in \mathcal{D}((b,t)) \\ (\perp, 0) & \text{otherwise} \end{cases}$$

while restrictions from a vertex a to an edge (a,b) are given by

$$\mathcal{D}\left((a,t) \subset ((a,t),(b,t))\right)(n_1, n_2, x_1, \ldots, x_n) = \begin{cases} (n_2, x_n) & \text{if } n_2 = a \text{ and } x_n \neq 0 \\ (n_2, x_1) & \text{if } n_2 \neq a \\ (\perp, 0) & \text{otherwise} \end{cases}$$

Proposition 22. *Every data payload sheaf contains an activation sheaf as a subsheaf when restricted to any timeslice.*

Proof. It is only necessary to match the definition of the stalks and restrictions within a timeslice, and to project onto the second component of each stalk over vertices and the first component over all other simplices. □

This means that the data payload sheaf incorporates the transmission structure described previously for activation sheaves, and more importantly that the within-timeslice restrictions for the data payload sheaf describe the relationship between the data on links and within the nodes.

For clarity, if $\mathcal{D}$ is a data payload sheaf on a time-dependent link complex X, the activation subsheaf at time t is written $A_t\mathcal{D}$. Therefore, there is a collection of surjections on stalks $A_t(a) : \mathcal{D}(a,t) \to A_t\mathcal{D}(a)$ that project out the appropriate components of the stalks as in the proof above. These surjections have the property that they are compatible with both sheaves, in that if $a \subset b$

$$A_t(b) \circ \mathcal{D}(a \subset b) = A_t\mathcal{D}(a \subset b) \circ A_t(a).$$

This is taken to be the description of a *sheaf morphism* $A_t : \mathcal{D} \to A_t\mathcal{D}$.

Proposition 23. *When restricted to a single node n, every data payload sheaf contains a 2-grouping sheaf taking values in the nodes adjacent to n as a subsheaf.*

Proof. By inspection of the restriction maps in the definitions. □

As a result, transmissions between timeslices are decoupled from one another. It is important to realize that this does *not* mean that the data payloads are decoupled. Instead, given a sequence of nodes that transmit at each timeslice – global sections of each activation sheaf in each timeslice – the data payload sheaf will describe the pathways for threading data through the network as the next proposition states.

Proposition 24. *Given a time-dependent link complex X and a data payload sheaf $\mathcal{D}$ and global sections $\{s_t\}$ for each activation subsheaf $A_t\mathcal{D}$, then*

1. *the restriction of each stalk $\mathcal{D}(a,t)$ to the collection of elements whose image through $A_t(a)$ is $s_t(a)$ yields a subsheaf $\mathcal{P}$ and*
2. *$\mathcal{P}$ is a sheaf taking values in the same category as the data payloads D, and*
3. *if D is a vector space, the dimension of the space of global sections of $\mathcal{P}$ is an upper bound on the network's throughput given the transmission pattern described by $\{s_t\}$.*

Proof.

1. The stalk of $\mathcal{P}$ over a vertex a at time t is given by $\mathcal{P}(a,t) = D^n$, the stalk over an edge e connecting two timeslices is $\mathcal{P}(e,t) = D^{n-1}$, and the stalk over any other simplex is merely D. The restrictions then all become projection maps between copies of D.
2. This is evident by construction, since the portion of each stalk of $\mathcal{D}$ that remains after fixing the $\{s_t\}$ consists of a product of copies of D.
3. Suppose it is possible that a packet x can transmitted from node a to node b with the activation pattern given by $\{s_t\}$. To prove the bound, we construct a global section of $\mathcal{P}$ that supports this particular data transfer. This means there is a sequence of nodes $a = n_1, n_2, \ldots, n_k = b$ through which the packet travels, and that n_i and n_{i+1} are adjacent nodes for each i. This means that those nodes must come active in that sequence, so there is a sequence $t_1 < \ldots < t_k$ in which $s_{t_i}(n_i) = n_i$ for all i and for which the data over the star of n_i at that time t_i is x for all i. This procedure defines a local section S of $\mathcal{P}$ on the vertices (n_i, t_i), which can easily be extended to be a section over all timeslices by assigning zero data to all simplices outside the region of influence of the active nodes. It remains to extend to the edges connecting timeslices, but this is merely a matter of assigning zeros to the yet-unassigned slots in the queues.

□

Beware that the upper bound given in Proposition 24 is not tight – there may be global sections that describe packets that do not reach their intended destination(s).

4 Next steps

The next logical theoretical steps for these analyses are the following:

1. *Analyzing packet flow through the network.* We aim for a theorem that shows that packets cannot overlap except if they're *waiting* in a queue at a node. They cannot overlap on higher dimensional faces within a timeslice.

2. *Protocol level jamming.* Showing how it can be possible to use regions of influence of particular routes to obstruct the network, possibly at particular times.

3. *Developing protocol models through a layered construction.* The resulting construction will likely leverage a filtration of sheaves or more generally a sequence of sheaf morphisms that describe increasing protocol complexity. In either of these settings, the abstract tool of spectral sequences may describe constraints imposed on higher layers from the lower ones.

Acknowledgement

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA) via Federal contract HR0011-15-C-0050. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

[1] A. Mosleh, N. Siu, C. Smidts, and C. Lui, "Model uncertainty: Its characterization and quantitfication." in *Proceedings of Workshop I in Advanced Topics in Risk and Reliability Analysis.*, no. NUREG/CP-0138, Washington, DC, 1994.

[2] L. KB, "Model uncertainty: theory and practical implications," *IEEE Trans Syst, Man, CybernPart A: Syst Hum*, vol. 26, pp. 340–348, 1996.